arXiv:1703.01680v3 [cs.LG] 19 Mar 2017

Multi-Objective Non-parametric Sequential Prediction

Guy Uziel
Technion – Israel Institute of Technology
Ran El-Yaniv
Technion – Israel Institute of Technology

March 21, 2017

Abstract

Online-learning research has mainly been focusing on minimizing one objective function. In many real-world applications, however, several objective functions have to be considered simultaneously. Recently, an algorithm for dealing with several objective functions in the i.i.d. case has been presented. In this paper, we extend the multi-objective framework to the case of stationary and ergodic processes, thus allowing dependencies among observations. We first identify an asymptomatic lower bound for any prediction strategy and then present an algorithm whose predictions achieve the optimal solution while fulfilling any continuous and convex constraining criterion.

1 Introduction

In the traditional online learning setting, and in particular in sequential prediction under uncertainty, the learner is evaluated by a single loss function that is not completely known at each iteration [9]. When dealing with multiple objectives, since it is impossible to simultaneously minimize all of the objectives, one objective is chosen as the main function to minimize, leaving the others to be bound by pre-defined thresholds. Methods for dealing with one objective function can be transformed to deal with several objective functions by giving each objective a pre-defined weight. The difficulty, however, lies in assigning an appropriate weight to each objective in order to keep the objectives below a given threshold. This approach is very problematic in real world applications, where the player is required to to satisfy certain constraints. For example, in online portfolio selection [16, 7], the player may want to maximize wealth while keeping the risk (i.e., variance) contained below a certain threshold. Another example is the Neyman-Pearson (NP) classification paradigm (see, e.g., [22]) (which extends the objective in classical binary classification) where the goal is to learn a classifier achieving low type II error whose type I error is kept below a given threshold.

Recently, [19] presented an algorithm for dealing with the case of one main objective and fully-known constraints. In a subsequent work, [20] proposed a framework for

dealing with multiple objectives in the stochastic i.i.d. case, where the learner does not have full information about the objective functions. They proved that if there exists a solution that minimizes the main objective function while keeping the other objectives below given thresholds, then their algorithm will converge to the optimal solution.

In this work, we study online prediction with multiple objectives but now consider the challenging general case where the unknown underlying process is stationary and ergodic, thus allowing observations to depend on each other arbitrarily. We consider a non-parametric approach, which has been applied successfully in various application domains. For example, in online portfolio selection, [14, 12], [13], and [17] proposed non-parametric online strategies that guarantee, under mild conditions, the best possible outcome. Another interesting example in this regard is the work on time-series prediction by [5], [11], and [6]. A common theme to all these results is that the asymptotically optimal strategies are constructed by combining the predictions of many simple experts. The algorithm presented in this paper utilizes as a sub-routine the Weak Aggregating Algorithm of [24], and [15] to handle multiple objectives. While we discuss here the case of only two objective functions, our theorems can be extended easily to any fixed number of functions.

Outline The paper is organized as follows: In Section 2, we define the multi-objective optimization framework under a jointly stationary and ergodic process. In Section 3, we identify an asymptotic lower-bound for any prediction strategy. In Section 4, we present Algorithm 1, which asymptotically achieves an optimal feasible solution.

2 Problem Formulation

We consider the following prediction game. Let $\mathcal{X} \triangleq [-D, D]^d \subset \mathbb{R}^d$ be a compact *observation space* where $D > 0$. At each round, $n = 1, 2, \ldots$, the player is required to make a prediction $y_n \in \mathcal{Y}$, where $\mathcal{Y} \subset \mathbb{R}^m$ is a compact and convex set, based on past observations, $X_1^{n-1} \triangleq (x_1, \ldots, x_{n-1})$ and, $x_i \in \mathcal{X}$ (X_1^0 is the empty observation). After making the prediction y_n, the observation x_n is revealed and the player suffers two losses, $u(y_n, x_n)$ and $c(y_n, x_n)$, where u and c are real-valued continuous functions and convex w.r.t. their first argument. We view the player's prediction strategy as a sequence $\mathcal{S} \triangleq \{S_n\}_{n=1}^{\infty}$ of forecasting functions $S_n : \mathcal{X}^{(n-1)} \to \mathcal{Y}$; that is, the player's prediction at round n is given by $S_n(X_1^{n-1})$. Throughout the paper we assume that $x_1, x_2, \ldots$ are realizations of random variables $X_1, X_2, \ldots$ such that the stochastic process $(X_n)_{-\infty}^{\infty}$ is jointly stationary and ergodic and $\mathbb{P}(X_i \in \mathcal{X}) = 1$. The player's goal is to play the game with a strategy that minimizes the average u-loss, $\frac{1}{N}\sum_{i=1}^{N} u(S(X_1^{i-1}), x_i)$, while keeping the average c-loss $\frac{1}{N}\sum_{i=1}^{N} c(S(X_1^{i-1}), x_i)$ bounded below a prescribed threshold γ. Formally, we define the following:

Definition 1 (γ-boundedness)**.** *A prediction strategy $\mathcal{S}$ will be called γ-bounded if*

$$\limsup_{N\to\infty} \left(\frac{1}{N} \sum_{i=1}^{N} c(S_i(X_1^{i-1}), X_i) \right) \leq \gamma$$

almost surely. The set of all γ-bounded strategies will be denoted $\mathcal{S}_\gamma$.

Definition 2 (γ-feasible process)**.** *We say that the stationary and ergodic process $\{X_i\}_{-\infty}^{\infty}$ is γ-feasible w.r.t. the functions u and c, if for $\mathbb{P}_\infty$, the regular conditional probability distribution of X_0 given $\mathcal{F}_\infty$ (the σ-algebra generated by the infinite past $X_{-1}, X_{-2}, \ldots$), and for a threshold $\gamma > 0$, if there exists some $y' \in \mathcal{Y}$ such that $\mathbb{E}_{\mathbb{P}_\infty}[c(y', X_0)] < \gamma$.*

If γ-feasibility holds, then we will denote by y_∞^* (y_∞^* is not necessarily unique) the solution to the following minimization problem:

$$\begin{aligned} \underset{y\in\mathcal{Y}}{\text{minimize}} \quad & \mathbb{E}_{\mathbb{P}_\infty}[u(y, X_0)] \\ \text{subject to} \quad & \mathbb{E}_{\mathbb{P}_\infty}[c(y, X_0)] \leq \gamma, \end{aligned} \tag{1}$$

(1) and we define the *γ-feasible optimal value* as

$$\mathcal{V}^* = \mathbb{E}\left[\mathbb{E}_{\mathbb{P}_\infty}[u(y_\infty^*, X_0)]\right] \quad a.s.$$

Note that problem (1) is a convex minimization problem over $\mathcal{Y}$, which in turn is a compact and convex subset of $\mathbb{R}^m$. Therefore, the problem is equivalent to finding the saddle point of the Lagrangian function [3], namely,

$$\min_{y\in\mathcal{Y}} \max_{\lambda\in\mathbb{R}^+} \mathcal{L}(y, \lambda),$$

where the Lagrangian is

$$\mathcal{L}(y, \lambda) \triangleq \left(\mathbb{E}_{\mathbb{P}_\infty}[u(y, X_0)] + \lambda\left(E_{\mathbb{P}_\infty}[c(y, X_0)] - \gamma\right)\right).$$

We denote the optimal dual by λ_∞^* and assume that λ_∞^* is unique. Moreover, we set a constant [1] $\lambda_{\max}$ such that $\lambda_{\max} > \lambda_\infty^*$, and set $\Lambda \triangleq [0, \lambda_{\max}]$. We also define the *instantaneous Lagrangian function* as

$$l(y, \lambda, x) \triangleq u(y, x) + \lambda\left(c(y, x) - \gamma\right). \tag{2}$$

In Brief, we are seeking a strategy $S \in \mathcal{S}_\gamma$ that is as good as any other γ-bounded strategy, in terms of the average u-loss, when the underlying process is γ-feasible. Such a strategy will be called γ-universal.

3 Optimallity of $\mathcal{V}^*$

In this section, we prove that the average u-loss of any γ-bounded prediction strategy cannot be smaller than $\mathcal{V}^*$, the γ-feasible optimal value. This result is a generalization of the well-known result of [1] regarding the best possible outcome under a single objective. Before stating and proving this optimallity result, we state one known lemma and state and prove two lemmas that will be used repeatedly in this paper. The first lemma is known as Breiman's generalized ergodic theorem. The second and the third lemmas concern the continuity of the saddle point w.r.t. the probability distribution.

[1] This can be done, for example, by imposing some regularity conditions on the constraint function (see, e.g., [20]).

Lemma 1 (Ergodicity, [8]). *Let $\mathbf{X} = \{X_i\}_{-\infty}^{\infty}$ be a stationary and ergodic process. For each positive integer i, let T_i denote the operator that shifts any sequence by i places to the left. Let $f_1, f_2, \ldots$ be a sequence of real-valued functions such that $\lim_{n\to\infty} f_n(\mathbf{X}) = f(\mathbf{X})$ almost surely, for some function f. Assume that $\mathbb{E}\sup_n |f_n(\mathbf{X})| < \infty$. Then,*

$$\lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} f_i(T^i \mathbf{X}) = \mathbb{E} f(\mathbf{X})$$

almost surely.

Lemma 2 (Continuity and Minimax). *Let $\mathcal{Y}, \Lambda, \mathcal{X}$ be compact real spaces. $l : \mathcal{Y} \times \Lambda \times \mathcal{X} \to \mathbb{R}$ be a continuous function. Denote by $\mathbb{P}(\mathcal{X})$ the space of all probability measures on $\mathcal{X}$ (equipped with the topology of weak-convergence). Then the following function $L^* : \mathbb{P}(\mathcal{X}) \to \mathbb{R}$ is continuous*

$$L^*(\mathbb{Q}) = \inf_{y\in\mathcal{Y}} \sup_{\lambda\in\Lambda} \mathbb{E}_{\mathbb{Q}} \left[l(y, \lambda, x) \right]. \tag{3}$$

Moreover, for any $\mathbb{Q} \in \mathbb{P}(\mathcal{X})$,

$$\inf_{y\in\mathcal{Y}} \sup_{\lambda\in\Lambda} \mathbb{E}_{\mathbb{Q}} \left[l(y, \lambda, x) \right] = \sup_{\lambda\in\Lambda} \inf_{y\in\mathcal{Y}} \mathbb{E}_{\mathbb{Q}} \left[l(y, \lambda, x) \right].$$

Proof. $\mathcal{Y}, \Lambda, \mathcal{X}$ are compact, implying that the function $l(y, \lambda, x)$ is bounded. Therefore, the function $L : \mathcal{Y} \times \Lambda \times \mathbb{P}(\mathcal{X}) \to \mathbb{R}$, defined as

$$L(y, \lambda, \mathbb{Q}) = \mathbb{E}_{\mathbb{Q}} \left[l(y, \lambda, x) \right], \tag{4}$$

is continuous. By applying Proposition 7.32 from [4], we have that $\sup_{\lambda\in\Lambda} \mathbb{E}_{\mathbb{Q}} \left[l(y, \lambda, X) \right]$ is continuous in $\mathbb{Q} \times \mathcal{Y}$. Again applying the same proposition, we get the desired result. The last part of the lemma follows directly from Fan's minimax theorem [10]. □

Lemma 3 (Continuity of the optimal selection). *Let $\mathcal{Y}, \Lambda, \mathcal{X}$ be compact real spaces, and let L be as defined in Equation (4). Then, there exist two measurable selection functions h^X,h^λ such that*

$$h^y(\mathbb{Q}) \in \arg\min_{y\in\mathcal{Y}} \left(\max_{\lambda\in\Lambda} L(y, \lambda, \mathbb{Q}) \right),$$
$$h^\lambda(\mathbb{Q}) \in \arg\max_{\lambda\in\Lambda} \left(\min_{y\in\mathcal{Y}} L(y, \lambda, \mathbb{Q}) \right)$$

for any $\mathbb{Q} \in \mathbb{P}(\mathcal{X})$. Moreover, let L^* be as defined in Equation (3). Then, the set

$$Gr(L^*) \triangleq \{(u^*, v^*, \mathbb{Q}) \mid u^* \in h^y(\mathbb{Q}), v^* \in h^\lambda(\mathbb{Q}), \mathbb{Q} \in \mathbb{P}(\mathcal{X})\},$$

is closed in $\mathcal{Y} \times \Lambda \times \mathbb{P}(\mathcal{X})$.

Proof. The first part of the proof follows immediately from the minimax measurable theorem of [21] due to the compactness of $\mathcal{Y}, \Lambda, \mathcal{X}$ and the properties of the loss function L. The proof of the second part is similar to the one presented in Theorem 3 of [2]. In order to show that $Gr(L^*)$ is closed, it is enough to show that if (i) $\mathbb{Q}_n \to \mathbb{Q}_\infty$ in $\mathbb{P}(\mathcal{X})$; (ii) $u_n \to u_\infty$ in $\mathcal{Y}$; (iii) $v_n \to v_\infty$ in Λ and (iv) $u_n \in h^y(\mathbb{Q}_n), v_n \in h^\lambda(\mathbb{Q}_n)$ for all n, then,

$$u_\infty \in h^y(\mathbb{Q}_\infty), v_\infty \in h^\lambda(\mathbb{Q}_\infty).$$

The function $L(y, \lambda, \mathbb{Q})$, as defined in Equation (4), is continuous. Therefore,

$$\lim_{n\to\infty} L(u_n, v_n, \mathbb{Q}_n) = L(u_\infty, v_\infty, \mathbb{Q}_\infty).$$

It remains to show that $u_\infty \in h^y(\mathbb{Q}_\infty)$ and $v_\infty \in h^\lambda(\mathbb{Q}_\infty)$. From the optimality of u_n and v_n, we obtain

$$L(u_\infty, v_\infty, \mathbb{Q}_\infty) = \lim_{n\to\infty} L(u_n, v_n, \mathbb{Q}_n) = \lim_{n\to\infty} L^*(\mathbb{Q}_n). \tag{5}$$

Finally, from the continuity of L^* (Lemma 2), we get

$$(5) = L^*(\lim_{n\to\infty} \mathbb{Q}_n) = L^*(\mathbb{Q}_\infty),$$

which gives the desired result. □

Corollary 1. *Under the conditions of Lemma 3. Define* $L_n(y, \lambda, \mathbb{Q}) = L(y, \lambda, \mathbb{Q}) + \frac{||y||^2 - ||\lambda||^2}{n}$ *and denote* $h^y_{L_n}(\mathbb{Q}_n), h^\lambda_{L_n}(\mathbb{Q}_n)$ *to be the measurable selection functions of* L_n. *If* $\mathbb{Q}_n \to \mathbb{Q}_\infty$ *weakly in* $\mathbb{P}(\mathcal{X})$ *and* $u_n \in h^y_{L_n}(\mathbb{Q}_n), v_n \in h^\lambda_{L_n}(\mathbb{Q}_n)$, *then*

$$L_n(u_n, v_n, \mathbb{Q}_n) \to L(u_\infty, v_\infty, \mathbb{Q}_\infty)$$

almost surely for $u_\infty \in h^y(\mathbb{Q}_\infty)$ *and* $v_\infty \in h^\lambda(\mathbb{Q}_\infty)$.

Proof. Denote $\hat{u}_n \in h^y(\mathbb{Q}_\infty)$ and $\hat{v}_n \in h^\lambda(\mathbb{Q}_\infty)$

$$\begin{aligned} &|L_n(u_n, v_n, \mathbb{Q}_n) - L(u_\infty, v_\infty, \mathbb{Q}_\infty)| \\ \leq\ &|L_n(u_n, v_n, \mathbb{Q}_n) - L(\hat{u}_n, \hat{v}_n, \mathbb{Q}_n)| + |L(\hat{u}_n, \hat{v}_n, \mathbb{Q}_n) - L(u_\infty, v_\infty, \mathbb{Q}_\infty)|. \end{aligned} \tag{6}$$

Note that for every n and for constant $E > 0$,

$$\begin{aligned} \min_{y\in\mathcal{Y}} \max_{\lambda\in\Lambda} L(y, \lambda, \mathbb{Q}) - \frac{||\lambda_{\max}||^2}{n} &\leq \min_{y\in\mathcal{Y}} \max_{\lambda\in\Lambda} L_n(y, \lambda, \mathbb{Q}) \\ = \min_{y\in\mathcal{Y}} \max_{\lambda\in\Lambda} &\left(\mathbb{E}_{\mathbb{Q}} \left[l(y, \lambda, X) \right] + \frac{||y||^2 - ||\lambda||^2}{n} \right) \\ \leq \min_{y\in\mathcal{Y}} \max_{\lambda\in\Lambda} &L(y, \lambda, \mathbb{Q}) + \frac{E}{n}. \end{aligned}$$

Thus, for some constant C, $|L_n(u_n, v_n, \mathbb{Q}_n) - L(u_\infty, v_\infty, \mathbb{Q}_\infty)| < \frac{C}{n}$ and from Lemma 3, the last summand also converges to 0 as n approaches ∞, we get the desired result, and clearly, if $h^y(\mathbb{Q}_\infty)$ and $h^\lambda(\mathbb{Q}_\infty)$ are singletons, then, the only accumulation point of $\{(v_n, u_n)\}_{n=1}^\infty$ is (v_∞, u_∞). □

The importance of Lemma 3 stems from the fact that it proves the continuity properties of the multi-valued correspondences $\mathbb{Q} \to h^y(\mathbb{Q})$ and $\mathbb{Q} \to h^\lambda(\mathbb{Q})$. This leads to the knowledge that if for the limiting distribution, $\mathbb{Q}_\infty$, the optimal set is a singleton, then $\mathbb{Q} \to h^y(\mathbb{Q})$ and $\mathbb{Q} \to h^\lambda(\mathbb{Q})$ are continuous in $\mathbb{Q}_\infty$. We are now ready to prove the optimality of $\mathcal{V}^*$.

Theorem 1 (Optimality of $\mathcal{V}^*$)**.** *Let $\{X_i\}_{-\infty}^{\infty}$ be a γ-feasible process. Then, for any strategy $\mathcal{S} \in \mathcal{S}_\gamma$, the following holds a.s.*

$$\liminf_{N\to\infty} \frac{1}{N} \sum_{i=1}^{N} u(S(X_1^{i-1}), X_i) \geq \mathcal{V}^*.$$

Proof. For any given strategy $\mathcal{S} \in \mathcal{S}_\gamma$, we will look at the following sequence:

$$\frac{1}{N} \sum_{i=1}^{N} l(S(X_1^{i-1}), \tilde{\lambda_i^*}, X_i). \tag{7}$$

where $\tilde{\lambda_i^*} \in h^\lambda(\mathbb{P}_{X_i|X_1^{i-1}})$ Observe that

$$(7) = \frac{1}{N} \sum_{i=1}^{N} \mathbb{E}\left[l(S(X_1^{i-1}), \tilde{\lambda_i^*}, X_i) \mid X_1^{i-1}\right] - \frac{1}{N} \sum_{i=1}^{N} (l(S(X_1^{i-1}), \tilde{\lambda_i^*}, X_i) - \mathbb{E}\left[l(S(X_1^{i-1}), \tilde{\lambda_i^*}, X) \mid X_1^{i-1}\right]).$$

Since $A_i = l(S(X_1^{i-1}), \tilde{\lambda_i^*}, X_i) - \mathbb{E}\left[l(S(X_1^{i-1}), \tilde{\lambda_i^*}, X_i) \mid X_1^{i-1}\right]$ is a martingale difference sequence, the last summand converges to 0 a.s., by the strong law of large numbers (see, e.g., [23]). Therefore,

$$\begin{aligned}\liminf_{N\to\infty} \frac{1}{N} \sum_{i=1}^{N} l(S(X_1^{i-1}), \tilde{\lambda_i^*}, X_i) &= \liminf_{N\to\infty} \frac{1}{N} \sum_{i=1}^{N} \mathbb{E}\left[l(S(X_1^{i-1}), \tilde{\lambda_i^*}, X_i) \mid X_1^{i-1}\right] \\ &\geq \liminf_{N\to\infty} \frac{1}{N} \sum_{i=1}^{N} \min_{y \in \mathcal{Y}()} \mathbb{E}\left[l(y, \tilde{\lambda_i^*}, X_i) \mid X_1^{i-1}\right], \end{aligned} \tag{8}$$

where the minimum is taken w.r.t. all the $\sigma(X_1^{i-1})$-measurable functions. Because the process is stationary, we get for $\hat{\lambda_i^*} \in h^\lambda(\mathbb{P}_{X_0|X_{1-i}^{-1}})$,

$$(8) = \liminf_{N\to\infty} \frac{1}{N} \sum_{i=1}^{N} \min_{y \in \mathcal{Y}()} \mathbb{E}\left[l(y, \hat{\lambda_i^*}, X_0) \mid X_{1-i}^{-1}\right] \tag{9}$$

$$= \liminf_{N\to\infty} \frac{1}{N} \sum_{i=1}^{N} L^*(\mathbb{P}_{X_0|X_{1-i}^{-1}}). \tag{10}$$

Using Levy's zero-one law, $\mathbb{P}_{X_0|X_{1-i}^{-1}} \to \mathbb{P}_\infty$ weakly as i approaches ∞ and from Lemma 2 we know that L^* is continuous. Therefore, we can apply Lemma 1 and get

that a.s.

$$(10) = \mathbb{E}\left[L^*(\mathbb{P}_\infty)\right] = \mathbb{E}\left[\mathbb{E}_{\mathbb{P}_\infty}\left[l\left(y^*_\infty, \lambda^*_\infty, X_0\right)\right]\right] = \mathbb{E}\left[\mathcal{L}\left(y^*_\infty, \lambda^*_\infty, X_0\right)\right]. \qquad (11)$$

Note also, that due to the complementary slackness condition of the optimal solution, i.e., $\lambda^*_\infty(\mathbb{E}_{\mathbb{P}_\infty}\left[c(y^*_\infty, X_0)\right] - \gamma) = 0$, we get

$$(11) = \mathbb{E}\left[\mathbb{E}_{\mathbb{P}_\infty}\left[u\left(y^*_\infty, X_0\right)\right]\right] = \mathcal{V}^*.$$

From the uniqueness of λ^*_∞, and using Lemma 3 $\hat{\lambda^*_i} \to \lambda^*_\infty$ as i approaches ∞. Moreover, since l is continuous on a compact set, l is also uniformly continuous. Therefore, for any given $\epsilon > 0$, there exists $\delta > 0$, such that if $|\lambda' - \lambda| < \delta$, then

$$|l(y, \lambda', x) - l(y, \lambda, x)| < \epsilon$$

for any $y \in \mathcal{Y}$ and $x \in \mathcal{X}$. Therefore, there exists i_0 such that if $i > i_0$ then $|l(y, \hat{\lambda^*_i}, x) - l(y, \lambda^*_\infty, x)| < \epsilon$ for any $y \in \mathcal{Y}$ and $x \in \mathcal{X}$. Thus,

$$\begin{aligned}
&\liminf_{N\to\infty} \frac{1}{N}\sum_{i=1}^{N} l(S(X_1^{i-1}), \lambda^*_\infty, X_i) - \liminf_{N\to\infty} \frac{1}{N}\sum_{i=1}^{N} l(S(X_1^{i-1}), \hat{\lambda^*_i}, X_i) \\
&= \liminf_{N\to\infty} \frac{1}{N}\sum_{i=1}^{N} l(S(X_1^{i-1}), \lambda^*_\infty, X_i) + \limsup_{N\to\infty} \frac{1}{N}\sum_{i=1}^{N} -l(S(X_1^{i-1}), \hat{\lambda^*_i}, X_i) \\
&\geq \liminf_{N\to\infty} \frac{1}{N}\sum_{i=1}^{N} l(S(X_1^{i-1}), \hat{\lambda^*_i}, X_i) - \frac{1}{N}\sum_{i=1}^{N} l(S(X_1^{i-1}), \lambda^*_\infty, X_i) \geq -\epsilon \ \ a.s.,
\end{aligned}$$

and since ϵ is arbitrary,

$$\liminf_{N\to\infty} \frac{1}{N}\sum_{i=1}^{N} l(S(X_1^{i-1}), \lambda^*_\infty, X_i) \geq \liminf_{N\to\infty} \frac{1}{N}\sum_{i=1}^{N} l(S(X_1^{i-1}), \hat{\lambda^*_i}, X_i).$$

Therefore we can conclude that

$$\liminf_{N\to\infty} \frac{1}{N}\sum_{i=1}^{N} l(S(X_1^{i-1}), \lambda^*_\infty, X_i) \geq \mathcal{V}^* \ \ a.s.$$

We finish the proof by noticing that since $\mathcal{S} \in \mathcal{S}_\gamma$, then by definition

$$\limsup_{N\to\infty} \frac{1}{N}\sum_{i=1}^{N} c(S(X_1^{i-1}), X_i) \leq \gamma \ \ a.s.$$

and since λ^*_∞ is non negative, we will get the desired result. □

The above lemma also provides the motivation to find the saddle point of the Lagrangian $\mathcal{L}$. Therefore, for the reminder of the paper we will use the loss function l as defined in Equation 2.

Algorithm 1 Minimax Histogram Based Aggregation (MHA)

Input: Countable set of experts $\{H_{k,h}\}$, $y_0 \in \mathcal{Y}$ $\lambda_0 \in \Lambda$, initial probability $\{\alpha_{k,h}\}$,
For $n = 0$ **to** ∞
Play y_n, λ_n.
Nature reveals x_n
Suffer loss $l(y_n, \lambda_n, x_n)$.
Update the cumulative loss of the experts

$$l_{y,n}^{k,h} \triangleq \sum_{i=0}^{n} l(y_{k,h}^{i}, \lambda_i, x_i) \qquad l_{\lambda,n}^{k,h} \triangleq \sum_{i=0}^{n} l(y_i, \lambda_{k,h}^{i}, x_i)$$

Update experts' weights

$$w_n^{y,(k,h)} \triangleq \alpha_{k,h} \exp\left(-\frac{1}{\sqrt{n}} l_{y,n}^{k,h}\right)$$
$$p_{n+1}^{y,(k,h)} \triangleq \frac{w_{n+1}^{y,(k,h)}}{\sum_{h=1}^{\infty}\sum_{k=1}^{\infty} w_{n+1}^{y,(k,h)}}$$

Update experts' weights $w_{n+1}^{\lambda,(k,h)}$

$$w_{n+1}^{\lambda,(k,h)} \triangleq \alpha_{k,h} \exp\left(\frac{1}{\sqrt{n}} l_{\lambda,n}^{k,h}\right)$$
$$p_{n+1}^{\lambda,(k,h)} = \frac{w_{n+1}^{\lambda,(k,h)}}{\sum_{h=1}^{\infty}\sum_{k=1}^{\infty} w_{n+1}^{\lambda,(k,h)}}$$

Choose y_{n+1} and λ_{n+1} as follows

$$y_{n+1} = \sum_{k,h} p_{n+1}^{y,(k,h)} y_{k,h}^{n+1} \quad \lambda_{n+1} = \sum_{k,h} p_{n+1}^{\lambda,(k,h)} \lambda_{k,h}^{n+1}$$

End For

4 Minimax Histogram Based Aggregation

We are now ready to present our algorithm *Minimax Histogram based Aggregation (MHA)* and prove that its predictions are as good as the best strategy. By Theorem 1 we can restate our goal: find a prediction strategy $\mathcal{S} \in \mathcal{S}_\gamma$ such that for any γ-feasible process $\{X_i\}_{-\infty}^{\infty}$ the following holds:

$$\lim_{N\to\infty} \frac{1}{N} \sum_{i=1}^{N} u(S(X_1^{i-1}), X_i) = \mathcal{V}^* \ \ a.s.$$

Such a strategy will be called γ-*universal*. We do so by maintaining a countable set of

experts $\{H_{k,h}\}$, where an expert $H_{k,l}$ will output a pair $(y^i_{k,l}, \lambda^i_{k,l}) \in \mathcal{Y} \times \Lambda$ at round i. Our algorithm outputs at round i a pair $(y_i, \lambda_i) \in \mathcal{Y} \times \Lambda$ where the sequence of predictions $y_1, y_2, \ldots$ tries to minimize the average loss $\frac{1}{N}\sum_{i=1}^{N} l(y, \lambda_i, x_i)$ and the sequence of predictions $\lambda_1, \lambda_2, \ldots$ tries to maximize the average loss $\frac{1}{N}\sum_{i=1}^{N} l(y_i, \lambda, x_i)$. Each of y_i and λ_i is the aggregation of predictions $y^i_{k,l}$ and $\lambda^i_{k,l}$, $k, l = 1, 2, \ldots$, respectively. In order to ensure that the performance of MHA will be as good as any other expert for both the y and the λ predictions, we apply the Weak Aggregating Algorithm of [24], and [15] twice simultaneously. In Theorem 2 we prove that there exists a countable set of experts whose selection of points converges to the optimal solution. Then, in Theorem 3 we prove that MHA applied on the experts defined in Theorem 2 generates a sequence of predictions that is γ-bounded and as good as any other strategy w.r.t. any γ-feasible process.

Theorem 2. *Assume that $\{X_i\}_{-\infty}^{\infty}$ is a γ-feasible process. Then, it is possible to construct a countable set of experts $\{H_{k,h}\}$ for which*

$$\lim_{k\to\infty}\lim_{h\to\infty}\lim_{n\to\infty}\frac{1}{N}\sum_{i=1}^{N} l(y^i_{k,h}, \lambda^i_{k,h}, X_i) = \mathcal{V}^* \;\; a.s.,$$

where $(y^i_{k,h}, \lambda^i_{k,h})$ are the predictions made by expert $H_{k,h}$ at round i.

Proof. We start by defining a countable set of experts $\{H_{k,h}\}$ as follow: For $h = 1, 2, \ldots$, let $P_h = \{A_{h,j} \mid j = 1, 2, ..., m_h\}$ be a sequence of finite partitions of $\mathcal{X}$ such that: (i) any cell of P_{h+1} is a subset of a cell of P_h for any h. Namely, P_{h+1} is a refinement of P_h; (ii) for a set A, if $diam(A) = \sup_{x,y\in A} ||x - y||$ denotes the diameter of A, then for any sphere B centered at the origin,

$$\lim_{h\to\infty}\max_{j:A_{h,j}\cap B\neq\emptyset} diam(A_{h,j}) = 0.$$

Define the corresponding quantizer $q_h(x) = j$, if $x \in A_{h,j}$. Thus, for any n and X_1^n, we define $Q_h(X_1^n)$ as the sequence $q_h(x_1), \ldots, q_h(x_n)$. For expert $H_{k,h}$, we define for $k > 0$, a k-long string of positive integers, denoted by w, the following set,

$$B^{w,(1,n-1)}_{k,h} \triangleq \{x_i \mid k < i < n, \;\; Q_h(X_{i-k}^{i-1}) = w\}.$$

We define also

$$h^y_{k,h}(X_1^{n-1}, w) \triangleq \arg\min_{y\in\mathcal{Y}}\left(\max_{\lambda\in\Lambda}\frac{1}{|B^{w,(1,n-1)}_{k,h}|}\sum_{x_i\in B^{w,(1,n-1)}_{k,h}} l_{k,l,n}(y, \lambda, x_i)\right)$$

$$h^\lambda_{k,h}(X_1^{n-1}, w) \triangleq \arg\max_{\lambda\in\Lambda}\left(\min_{y\in\mathcal{Y}}\frac{1}{|B^{w,(1,n-1)}_{k,h}|}\sum_{x_i\in B^{w,(1,n-1)}_{k,h}} l_{k,l,n}(y, \lambda, x_i)\right)$$

for

$$l_{k,h,n}(y, \lambda, x) \triangleq l(y, \lambda, x) + \left(||y||^2 - ||\lambda||^2\right)\left(\frac{1}{n} + \frac{1}{h} + \frac{1}{k}\right)$$

and we will set $h^y_{k,h}(X_1^{n-1}, w) = y_0$ and $h^\lambda_{k,h}(X_1^{n-1}, w) = \lambda_0$ for arbitrary $(y_0, \lambda_0) \in \mathcal{Y} \times \Lambda$ if $B^{w,(1,n-1)}_{k,h}$ is empty. Using the above, we define the predictions of $H_{k,h}$ to be:

$$H^y_{k,h}(X_1^{n-1}) = h^y_{k,h}(X_1^{n-1}, Q(X_{n-k}^{n-1})),\ n = 1, 2, 3....$$
$$H^\lambda_{k,h}(X_1^{n-1}) = h^\lambda_{k,h}(X_1^{n-1}, Q(X_{n-k}^{n-1})),\ n = 1, 2, 3....$$

We will add two experts: $H_{0,0}$ whose predictions are always $(y_0, \lambda_{\max})$ and $H_{-1,-1}$ whose predictions are always $(y_0, 0)$.

Fixing $k, h > 0$ and w, we will define a (random) measure $\mathbb{P}^{(k.h)}_{j,w}$ that is the measure concentrated on the set $B^{w,(0,1-j)}_{k,h}$, defined by

$$\mathbb{P}^{(k,h)}_{j,w}(A) = \frac{\sum_{X_i \in B^{w,(0,1-j)}_{k,h}} 1_A(X_i)}{|B^{w,(0,1-j)}_{k,h}|},$$

where 1_A denotes the indicator function of the set $A \subset \mathcal{X}$. If the above set $B^w_{k,h}$ is empty, then let $P^{(k,h)}_{j,w}(A) = \delta(x')$ be the probability measure concentrated on arbitrary vector $x' \in \mathcal{X}$.

In other words, $\mathbb{P}^{(k.h)}_{j,w}(A)$ is the relative frequency of the the vectors among $X_{1-j+k}, \ldots, X_0$ that fall in the set A. Applying Lemma 1 twice, it is straightforward to prove that for all w, w.p. 1

$$\mathbb{P}^{(k,h)}_{j,w} \to \begin{cases} \mathbb{P}_{X_0|G_l(X_{-k}^{-1})=w} & \mathbb{P}(G_l(X_{-k}^{-1}) = w) > 0 \\ \delta(x') & \textit{otherwise} \end{cases}$$

weakly as $j \to \infty$, where $\mathbb{P}_{X_0|G_l(X_{-k}^{-1})=w}$ denotes the distribution of the vector X_0 conditioned on the event $G_l(X_{-k}^{-1}) = w$. To see this, let f be a bounded continuous function. Then,

$$\int f(x)\mathbb{P}^{(k,h)}_{j,w}(dx) = \frac{\frac{1}{|1-j+k|}\sum_{X_i \in B^{w,(0,1-j)}_{k,h}} f(X_i)}{\frac{1}{|1-j+k|}|B^{w,(0,1-j)}_{k,h}|}$$
$$\to \frac{\mathbb{E}\left[f(X_0)1_{G_l(X_{-k}^{-1})=w}(X_0)\right]}{\mathbb{P}(G_l(X_{-k}^{-1}) = w)} = \mathbb{E}\left[f(X_0) \mid G_l(X_{-k}^{-1}) = w\right],$$

and in case $\mathbb{P}(||X_{-k}^{-1} - s|| \le c/l) = 0$, then w.p. 1, $\mathbb{P}^{(k,h)}_{j,w}$ is concentrated on x' for all j. We will denote the limit distribution of $\mathbb{P}^{(k,h)}_{j,w}$ by $\mathbb{P}^{*(k,h)}_{w}$.

By definition, $\left(h^y_{k,h}(X_{1-n}^{-1}, w), h^\lambda_{k,h}(X_{1-n}^{-1}, w)\right)$ is the minimax of $l_{n,k,h}$ w.r.t. $\mathbb{P}^{(k,h)}_{j,w}$. The sequence of functions $l_{n,k,h}$ converges uniformly as n approaches ∞ to

$$l_{k,h}(y, \lambda, x) = l(y, \lambda, x) + \left(||y||^2 - ||\lambda||^2\right)\left(\frac{1}{h} + \frac{1}{k}\right).$$

Note also that for any fixed $\mathbb{Q}$, $L_{k,h}(y,\lambda,\mathbb{Q}) = \mathbb{E}_{\mathbb{Q}}\left[l_{k,h}(y,\lambda,X)\right]$ is strictly convex in y and strictly concave in λ, and therefore, has a unique saddle-point (see, e.g., [18]). Therefore, since w is arbitrary, and following a Corollary 1 of Lemma 3, we get that a.s.

$$y^n_{k,h} \to y^*_{k,h} \qquad \lambda^n_{k,h} \to \lambda^*_{k,h},$$

where $\left(y^*_{k,h}, \lambda^*_{k,h}\right)$ is the minimax of $L_{k,h}$ w.r.t. $\mathbb{P}^{*(k,h)}_{X^{-1}_{-k}}$. Thus, we can apply Lemma 1 and conclude that as N approaches ∞,

$$\frac{1}{N}\sum_{i=1}^{N} l(y^i_{k,h}, \lambda^i_{k,h}, X_i) \to \mathbb{E}\left[l(y^*_{k,h}, \lambda^*_{k,h}, X_0)\right].$$

a.s.. We now evaluate

$$\lim_{h\to\infty} \mathbb{E}\left[l(y^*_{k,h}, \lambda^*_{k,h}, X_0)\right].$$

Using the properties of the partition P_h (see, e.g., [11, 13]), we get that

$$\mathbb{P}^{*(k,h)}_{X^{-1}_{-k}} \to \mathbb{P}_{\left\{X_0|X^{-1}_{-k}\right\}}$$

weakly as $h \to \infty$. Moreover, the sequence of functions $l_{k,h}$ converges uniformly as h approaches ∞

$$l_k(y,\lambda,x) = l(y,\lambda,x) + \frac{||y||^2 - ||\lambda||^2}{k}.$$

Note also, that for any fixed $\mathbb{Q}$, $L_k(y,\lambda,\mathbb{Q}) = \mathbb{E}_{\mathbb{Q}}\left[l_k(y,\lambda,X)\right]$ is strictly convex-concave, and therefore, has a unique saddle point. Accordingly, by applying Corollary 1 again, we get that a.s.

$$y^*_{k,h} \to y^*_k \qquad \lambda^*_{k,h} \to \lambda^*_k,$$

where (y^*_k, λ^*_k) is the minimax of L_k w.r.t. $\mathbb{P}_{\left\{X_0|X^{-1}_{-k}\right\}}$. Therefore, as h approaches ∞,

$$l(y^*_{k,h}, \lambda^*_{k,h}, X_0) \to l\left(y^*_k, \lambda^*_k, X_0\right)$$

a.s.. Thus, by Lebesgue's dominated convergence,

$$\lim_{h\to\infty} \mathbb{E}\left[l(y^*_{k,h}, \lambda^*_{k,h}, X_0)\right] = \mathbb{E}\left[l\left(y^*_k, \lambda^*_k, X_0\right)\right].$$

Notice that for any $\mathbb{Q} \in \mathbb{P}(\mathcal{X})$, the distance between the saddle point of L_k w.r.t. $\mathbb{Q}$ and the the saddle point of L w.r.t. $\mathbb{Q}$ converges to 0 as k approaches ∞. To see this, notice that

$$\begin{aligned}\min_{y\in\mathcal{Y}}\max_{\lambda\in\Lambda} L(y,\lambda,\mathbb{Q}) - \frac{||\lambda_{\max}||^2}{k} &\le \min_{y\in\mathcal{Y}}\max_{\lambda\in\Lambda} L_k(y,\lambda,\mathbb{Q}) \\ &\le \min_{y\in\mathcal{Y}}\max_{\lambda\in\Lambda} L(y,\lambda,\mathbb{Q}) + \frac{E}{k}\end{aligned}$$

for some constant E, since $\mathcal{Y}$ is bounded. The last part in our proof will be to show that if $(\hat{y_k^*}, \hat{\lambda_k^*})$ is the minimax of L w.r.t. $\mathbb{P}_{\left\{X_0|X_{-k}^{-1}\right\}}$, then as k approaches ∞, $\mathbb{E}\left[l\left(\hat{y_k^*}, \hat{\lambda_k^*}, X_0\right)\right]$ will converge a.s. to $\mathcal{V}^*$ and so $\mathbb{E}\left[l\left(y_k^*, \lambda_k^*, X_0\right)\right]$.

To show this, we will use the sub-martingale convergence theorem twice. First, we define Z_k as

$$Z_k \triangleq \min_{y \in \mathcal{Y}()} \mathbb{E}\left[\max_{\lambda \in \Lambda()} \mathbb{E}\left[l\left(y, \lambda, X_0\right) \mid X_{-\infty}^{-1}\right] \mid X_{-k}^{-1}\right],$$

where the minimum is taken w.r.t. all $\sigma(X_{-k}^{-1})$-measurable strategies and the maximum is taken w.r.t. all $\sigma(X_{-\infty}^{-1})$-measurable strategies. Notice that Z_k is a super-martingale. We can see this by using the tower property of conditional expectations,

$$\mathbb{E}[Z_{k+1} \mid X_{-k}^{-1}] = \mathbb{E}\left[\mathbb{E}\left[Z_{k+1} \mid X_{-k-1}^{-1}\right] \mid X_{-k}^{-1}\right]$$

and since Z_{k+1} is the optimal choice in $\mathcal{Y}$ w.r.t. to X_{-k-1}^{-1},

$$\leq \mathbb{E}\left[\mathbb{E}[Z_k \mid X_{-k-1}^{-1}] \mid X_{-k}^{-1}\right] = \mathbb{E}[Z_k \mid X_{-k}^{-1}] = Z_k.$$

Note also that $\mathbb{E}[Z_k]$ is uniformly bounded. Therefore, we can apply the super-martingale convergence theorem and get that $Z_k \to Z_\infty$ a.s., where,

$$Z_\infty = \mathbb{E}\left[l(y_\infty^*, \lambda_\infty^*, X_0) \mid X_{-\infty}^{-1}\right] = \mathcal{V}^*,$$

and by using Lebesgue's dominated convergence theorem, also $\mathbb{E}[Z_k] \to \mathbb{E}[Z_\infty] = \mathcal{V}^*$. Using the same arguments, Z_k', defined as

$$Z_k' \triangleq \max_{\lambda \in \Lambda()} \mathbb{E}\left[\min_{y \in \mathcal{Y}()} \mathbb{E}\left[l\left(y, \lambda, X_0\right) \mid X_{-\infty}^{-1}\right] \mid X_{-k}^{-1}\right],$$

where the maximum is taken w.r.t. all $\sigma(X_{-k}^{-1})$-measurable strategies and the minimum is taken w.r.t. all $\sigma(X_{-\infty}^{-1})$-measurable strategies, is a sub-martingale that also converges a.s. to Z_∞ and thus $\mathbb{E}[Z_k'] \to \mathbb{E}[Z_\infty] = \mathcal{V}^*$.

We conclude the proof by noticing that the following relation holds for any k,

$$\begin{aligned}
\mathbb{E}[Z_k'] &= \mathbb{E}\left[\max_{\lambda \in \Lambda()} \mathbb{E}\left[\min_{y \in \mathcal{Y}()} \mathbb{E}\left[l\left(y, \lambda, X_0\right) \mid X_{-\infty}^{-1}\right] \mid X_{-k}^{-1}\right]\right] \\
&\leq \mathbb{E}\left[\max_{\lambda \in \Lambda()} \mathbb{E}\left[\mathbb{E}\left[l\left(\hat{y_k^*}, \lambda, X_0\right) \mid X_{-\infty}^{-1}\right] \mid X_{-k}^{-1}\right]\right] \\
&= \mathbb{E}\left[\max_{\lambda \in \Lambda()} \mathbb{E}\left[l\left(\hat{y_k^*}, \lambda, X_0\right) \mid X_{-k}^{-1}\right]\right] = \mathbb{E}\left[l\left(\hat{y_k^*}, \hat{\lambda_k^*}, X_0\right)\right],
\end{aligned}$$

and using similar arguments we can show that also

$$\mathbb{E}\left[l\left(\hat{y_k^*}, \hat{\lambda_k^*}, X_0\right)\right] \leq \mathbb{E}[Z_k],$$

and since both $\mathbb{E}[Z_k]$ and $\mathbb{E}[Z_k']$ converge to $\mathcal{V}^*$, we get the desired result.

□

Before stating the main theorem regarding MHA, we now state and prove the following lemma, which is used in the proof of the main result regarding MHA.

Lemma 4. *Let $\{H_{k,h}\}$ be a countable set of experts as defined in the proof of Theorem 2. Then, the following relation holds a.s.:*

$$\inf_{k,h} \limsup_{n\to\infty} \frac{1}{N} \sum_{i=1}^{N} l\left(y_{k,h}^{i}, \lambda_i, X_i\right) \leq \mathcal{V}^*$$
$$\leq \sup_{k,h} \liminf_{n\to\infty} \frac{1}{N} \sum_{i=1}^{N} l\left(y_i, \lambda_{k,h}^{i}, X_i\right),$$

where (y_i, λ_i) are the predictions of MHA when applied on $\{H_{k,h}\}$.

Proof. Set

$$f(y, \mathbb{Q}) \triangleq \max_{\lambda\in\Lambda} \mathbb{E}_{\mathbb{Q}}\left[l\left(y, \lambda, X_0\right)\right].$$

We will start from the LHS,

$$\inf_{k,h} \limsup_{n\to\infty} \frac{1}{N} \sum_{i=1}^{N} l\left(y_{k,h}^{i}, \lambda_i, X_i\right), \tag{12}$$

and similarly to Lemma 1, by using the strong law of large numbers we can write

$$(12) = \inf_{k,h} \limsup_{n\to\infty} \frac{1}{N} \sum_{i=1}^{N} \mathbb{E}\left[l\left(y_{k,h}^{i}, \lambda_i, X_0\right) \mid X_{1-i}^{-1}\right]$$
$$\leq \inf_{k,h} \limsup_{n\to\infty} \frac{1}{N} \sum_{i=1}^{N} f(y_{k,h}^{i}, \mathbb{P}_{X_0|X_{1-i}^{-1}}) \ \ a.s. \tag{13}$$

For fixed $k, h > 0$, from the proof of Theorem (2), $y_{k,h}^{i} \to y_{k,h}^{*}$ a.s. as i approaches ∞, and from Levy's zero-one law also $\mathbb{P}_{X_0|X_{1-i}^{-1}} \to \mathbb{P}_{\infty}$ weakly. From Lemma 2 we know that f is continuous, therefore, we can apply Lemma 1 and get that

$$(13) = \inf_{k,h} \mathbb{E}\left[\mathbb{E}\left[f(y_{k,h}^{*}, \mathbb{P}_{\infty})\right]\right] \leq \lim_{k\to\infty} \lim_{l\to\infty} \mathbb{E}\left[f(y_{k,h}^{*}, \mathbb{P}_{\infty})\right]. \tag{14}$$

From the uniqueness of the saddle point and from the proof of Theorem (2), for fiked $k > 0$,

$$\lim_{h\to\infty} y_{k,h}^{*} \to y_{k}^{*}$$

a.s.. Thus, from the continuity of f we get that

$$\lim_{h\to\infty} f(y_{k,h}^{*}, \mathbb{P}_{\infty}) \to f(y_{k}^{*}, \mathbb{P}_{\infty})$$

and again by Lebesgue's dominated convergence,

$$(14) = \lim_{k\to\infty} \mathbb{E}\left[f(y_{k}^{*}, \mathbb{P}_{\infty})\right] = \lim_{k\to\infty} \mathbb{E}\left[\max_{\lambda\in\Lambda} \mathbb{E}_{\mathbb{P}_{\infty}}\left[l\left(y_{k}^{*}, \lambda, X_0\right)\right]\right]. \tag{15}$$

Now, from Theorem 2 we know that every accumulation point of the sequence $\{y_k^*\}$ is in the optimal set

$$\arg\min_{y\in\mathcal{Y}}\left(\max_{\lambda\in\Lambda}\mathbb{E}_{\mathbb{P}_\infty}\left[l\left(y,\lambda,X_0\right)\right]\right).$$

Therefore a.s.

$$\lim_{k\to\infty}\max_{\lambda\in\Lambda}\mathbb{E}_{\mathbb{P}_\infty}\left[l\left(y_k^*,\lambda,X_0\right)\right]\to\mathbb{E}_{\mathbb{P}_\infty}\left[l\left(y_\infty^*,\lambda_\infty^*,X_0\right)\right],$$

and using Lebesgue's dominated convergence,

$$(15)=\mathbb{E}\left[\mathbb{E}_{\mathbb{P}_\infty}\left[l\left(y_\infty^*,\lambda_\infty^*,X_0\right)\right]\right]=\mathcal{V}^*.$$

Using similar arguments, we can show the second part of the lemma.

□

We are now ready to state and prove the optimality of MHA.

Theorem 3 (Optimality of MHA)**.** *Let (y_i,λ_i) be the predictions generated by MHA when applied on $\{H_{k,h}\}$ as defined in the proof of Theorem 2. Then, for any γ-feasible process $\{X_i\}_{-\infty}^{\infty}$: MHA is a γ-bounded and γ-universal strategy.*

Proof. We first show that

$$\lim_{N\to\infty}\frac{1}{N}\sum_{i=1}^{N}l(y_i,\lambda_i,X_i)=\mathcal{V}^*\ \ a.s. \tag{16}$$

Applying Lemma 5 in [15], we know that the x updates guarantee that for every expert $H_{k,h}$,

$$\frac{1}{N}\sum_{i=1}^{N}l(y_i,\lambda_i,x_i)\le\frac{1}{N}\sum_{i=1}^{N}l(y_{k,h}^i,\lambda_i,x_i)+\frac{C_{k,h}}{\sqrt{N}} \tag{17}$$

$$\frac{1}{N}\sum_{i=1}^{N}l(y_i,\lambda_i,x_i)\ge\frac{1}{N}\sum_{i=1}^{N}l(y_i,\lambda_{k,h}^i,x_i)-\frac{C'_{k,h}}{\sqrt{N}}, \tag{18}$$

where $C_{k,h},C'_{k,h}>0$ are some constants independent of N. In particular, using Equation (17),

$$\frac{1}{N}\sum_{i=1}^{N}l(y_i,\lambda_i,x_i)\le\inf_{k,h}\left(\frac{1}{N}\sum_{i=1}^{N}l(y_{k,h}^i,\lambda_i,x_i)+\frac{C_{k,h}}{\sqrt{N}}\right).$$

Therefore, we get

$$\begin{aligned}
&\limsup_{N\to\infty} \frac{1}{N}\sum_{i=1}^{N} l(y_i,\lambda_i,x_i) \\
&\leq \limsup_{N\to\infty} \inf_{k,h} \left(\frac{1}{N}\sum_{i=1}^{N} l(y^i_{k,h},\lambda_i,x_i) + \frac{C_{k,h}}{\sqrt{N}} \right) \\
&\leq \inf_{k,h} \limsup_{N\to\infty} \left(\frac{1}{N}\sum_{i=1}^{N} l(y^i_{k,h},\lambda_i,x_i) + \frac{C_{k,h}}{\sqrt{N}} \right) \\
&\leq \inf_{k,h} \limsup_{N\to\infty} \left(\frac{1}{N}\sum_{i=1}^{N} l(y^i_{k,h},\lambda_i,x_i) \right), \qquad (19)
\end{aligned}$$

where in the last inequality we used the fact that $\limsup$ is sub-additive. Using Lemma (4), we get that

$$\begin{aligned}
(19) &\leq \mathcal{V}^* \\
&\leq \sup_{k,h} \liminf_{n\to\infty} \frac{1}{N}\sum_{i=1}^{N} l\left(y_i,\lambda^i_{k,h},X_i\right). \qquad (20)
\end{aligned}$$

Using similar arguments and using Equation (18) we can show that

$$(20) \leq \liminf_{N\to\infty} \frac{1}{N}\sum_{i=1}^{N} l(y_i,\lambda_i,x_i).$$

Summarizing, we have

$$\limsup_{N\to\infty} \frac{1}{N}\sum_{i=1}^{N} l(y_i,\lambda_i,x_i) \leq \mathcal{V}^* \leq \liminf_{N\to\infty} \frac{1}{N}\sum_{i=1}^{N} l(y_i,\lambda_i,x_i).$$

Therefore, we can conclude that a.s.

$$\lim_{N\to\infty} \frac{1}{N}\sum_{i=1}^{N} l(y_i,\lambda_i,X_i) = \mathcal{V}^*.$$

To show that MHA is indeed a γ-bounded strategy and to shorten the notation, we will denote

$$g(y,\lambda,x) \triangleq \lambda(c(y,x)-\gamma).$$

First, from Equation (18) applied on the expert $H_{0,0}$, we get that:

$$\limsup_{N\to\infty} \frac{1}{N}\sum_{i=1}^{N} g(y_i,\lambda_{\max},x) \leq \limsup_{N\to\infty} \frac{1}{N}\sum_{i=1}^{N} g(y_i,\lambda_i,x). \qquad (21)$$

Moreover, since l is uniformly continuous, for any given $\epsilon > 0$, there exists $\delta > 0$, such that if $|\lambda' - \lambda| < \delta$, then

$$|l(y, \lambda', x) - l(y, \lambda, x)| < \epsilon$$

for any $y \in \mathcal{Y}$ and $x \in \mathcal{X}$. We also know that

$$\lim_{k\to\infty} \lim_{h\to\infty} \lim_{i\to\infty} \lambda^i_{k,h} = \lambda^*_\infty.$$

Therefore, there exist k_0, h_0, i_0 such that $|\lambda^i_{k_0,h_0} - \lambda^*_\infty| < \delta$ for any $i > i_0$. Since $\lim_{k\to\infty} \lambda^*_k = \lambda^*_\infty$ there exists k_0 such that $|\lambda^*_{k_0} - \lambda^*_\infty| < \frac{\delta}{3}$. Note that $\lim_{h\to\infty} \lambda^*_{k_0,h} = \lambda^*_{k_0}$, so there exists h_0 such that $|\lambda^*_{k_0,h_0} - \lambda^*_{k_0}| < \frac{\delta}{3}$. Finally, since $\lim_{i\to\infty} \lambda^i_{k_0,l_0} = \lambda^*_{k_0,l_0}$, there exists i_0 such that if $i > i_0$, then $|\lambda^i_{k_0,l_0} - \lambda^*_{k_0,l_0}| < \frac{\delta}{3}$. Combining all the above, we get that for k_0, h_0, i_0 if $i > i_0$, then

$$|\lambda^i_{k_0,h_0} - \lambda^*_\infty| < |\lambda^i_{k_0,h_0} - \lambda^*_{k_0,h_0}| + |\lambda^i_{k_0,h_0} - \lambda^*_{k_0}| + |\lambda^*_{k_0} - \lambda^*_\infty| < \delta.$$

Therefore,

$$\begin{aligned}
&\limsup_{N\to\infty} \left(\frac{1}{N} \sum_{i=1}^N l(y_i, \lambda^*_\infty, x_i) - \frac{1}{N} \sum_{i=1}^N l(y_i, \lambda_i, x_i) \right) \leq \\
&\limsup_{N\to\infty} \left(\frac{1}{N} \sum_{i=1}^N l(y_i, \lambda^*_\infty, x_i) - \frac{1}{N} \sum_{i=1}^N l(y_i, \lambda^i_{k_0,h_0}, x_i) \right) + \\
&\limsup_{N\to\infty} \left(\frac{1}{N} \sum_{i=1}^N l(y_i, \lambda^i_{k_0,h_0}, x_i) - \frac{1}{N} \sum_{i=1}^N l(y_i, \lambda_i, x_i) \right) \qquad (22)
\end{aligned}$$

From the uniform continuity we also learn that the first summand is bounded above by ϵ, and from Equation (18), we get that the last summand is bounded above by 0. Thus,

$$(22) \leq \epsilon,$$

and since ϵ is arbitrary, we get that

$$\limsup_{N\to\infty} \left(\frac{1}{N} \sum_{i=1}^N l(y_i, \lambda^*_\infty, x_i) - \frac{1}{N} \sum_{i=1}^N l(y_i, \lambda_i, x_i) \right) \leq 0.$$

Thus,

$$\limsup_{N\to\infty} \frac{1}{N} \sum_{i=1}^N l(y_i, \lambda^*_\infty, X_i) \leq \mathcal{V}^*,$$

and from Theorem 1 we can conclude that

$$\lim_{N\to\infty} \frac{1}{N} \sum_{i=1}^N l(y_i, \lambda^*_\infty, X_i) = \mathcal{V}^*.$$

Therefore, we can deduce that

$$\begin{aligned}
&\limsup_{N\to\infty} \frac{1}{N}\sum_{i=1}^{N} g(y_i,\lambda_i,x_i) - \limsup_{N\to\infty} \frac{1}{N}\sum_{i=1}^{N} g(y_i,\lambda^*_\infty,x_i) = \\
&\limsup_{N\to\infty} \frac{1}{N}\sum_{i=1}^{N} g(y_i,\lambda_i,x_i) + \liminf_{N\to\infty} \frac{1}{N}\sum_{i=1}^{N} -g(y_i,\lambda^*_\infty,x_i) \\
&\leq \limsup_{N\to\infty} \frac{1}{N}\sum_{i=1}^{N} g(y_i,\lambda_i,x_i) - \frac{1}{N}\sum_{i=1}^{N} g(y_i,\lambda^*_\infty,x_i) \\
&= \limsup_{N\to\infty} \frac{1}{N}\sum_{i=1}^{N} l(y_i,\lambda_i,x_i) - \frac{1}{N}\sum_{i=1}^{N} l(y_i,\lambda^*_\infty,x_i) = 0,
\end{aligned}$$

which results in

$$\limsup_{N\to\infty} \frac{1}{N}\sum_{i=1}^{N} g(y_i,\lambda_i,x_i) \leq \limsup_{N\to\infty} \frac{1}{N}\sum_{i=1}^{N} g(y_i,\lambda^*_\infty,x_i).$$

Combining the above with Equation (21), we get that

$$\begin{aligned}
&\limsup_{N\to\infty} \frac{1}{N}\sum_{i=1}^{N} g(y_i,\lambda_{\max},x_i) \\
&\leq \limsup_{N\to\infty} \frac{1}{N}\sum_{i=1}^{N} g(y_i,\lambda^*_\infty,x_i).
\end{aligned}$$

Since $0 \leq \lambda^*_\infty < \lambda_{\max}$, we get that MHA is γ-bounded. This also implies that

$$\limsup_{N\to\infty} \frac{1}{N}\sum_{i=1}^{N} \lambda_i(c(y_i,x_i)-\gamma) \leq 0.$$

Now, if we apply Equation (18) on the expert $H_{-1,-1}$, we get that

$$\liminf_{N\to\infty} \frac{1}{N}\sum_{i=1}^{N} \lambda_i(c(y_i,x_i)-\gamma) \geq 0.$$

Thus,

$$\lim_{N\to\infty} \frac{1}{N}\sum_{i=1}^{N} \lambda_i(c(y_i,x_i)-\gamma) = 0,$$

and using Equation (16), we get that MHA is also γ-universal.

□

5 Concluding Remarks

In this paper, we introduced the Minimax Histogram Aggregation (MHA) algorithm for multiple-objective sequential prediction. We considered the general setting where the unknown underlying process is stationary and ergodic., and given that the underlying process is γ-feasible, we extended the well-known result of [1] regarding the asymptotic lower bound of prediction with a single objective, to the case of multi-objectives. We proved that MHA is a γ-bounded strategy whose predictions also converge to the optimal solution in hindsight.

In the proofs of the theorems and lemmas above, we used the fact that the initial weights of the experts, $\alpha_{k,h}$, are strictly positive thus implying a countably infinite expert set. In practice, however, one cannot maintain an infinite set of experts. Therefore, it is customary to apply such algorithms with a finite number of experts (see [14, 12, 13, 17]). Despite the fact that in the proof we assumed that the observation set $\mathcal{X}$ is known a priori, the algorithm can also be applied in the case that $\mathcal{X}$ is unknown by applying the doubling trick. For a further discussion on this point, see [11]. In our proofs, we relied on the compactness of the set $\mathcal{X}$. It will be interesting to see whether the universality of MHA can be sustained under unbounded processes as well. A very interesting open question would be to identify conditions allowing for finite sample bounds when predicting with multiple objectives.

References

[1] P.H. Algoet. The strong law of large numbers for sequential decisions under uncertainty. *IEEE Transactions on Information Theory*, 40(3):609–633, 1994.

[2] P.H. Algoet and T.M. Cover. Asymptotic optimality and asymptotic equipartition properties of log-optimum investment. *The Annals of Probability*, pages 876–898, 1988.

[3] A. Ben-Tal and A. Nemirovsky. Optimization iii. *Lecture Notes*, 2012.

[4] D. Bertsekas and S. Shreve. *Stochastic optimal control: The discrete time case*, volume 23. Academic Press New York, 1978.

[5] G. Biau, K. Bleakley, L. Györfi, and G. Ottucsák. Nonparametric sequential prediction of time series. *Journal of Nonparametric Statistics*, 22(3):297–317, 2010.

[6] G. Biau and B. Patra. Sequential quantile prediction of time series. *IEEE Transactions on Information Theory*, 57(3):1664–1674, 2011.

[7] A. Borodin and R. El-Yaniv. *Online Computation and Competitive Analysis*. Cambridge University Press, 2005.

[8] L. Breiman. The individual ergodic theorem of information theory. *The Annals of Mathematical Statistics*, 28(3):809–811, 1957.

[9] N. Cesa-Bianchi and G. Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.

[10] K. Fan. Minimax theorems. *Proceedings of the National Academy of Sciences*, 39(1):42–47, 1953.

[11] L. Györfi and G. Lugosi. Strategies for sequential prediction of stationary time series. In *Modeling uncertainty*, pages 225–248. Springer, 2005.

[12] L. Györfi, G. Lugosi, and F. Udina. Nonparametric kernel-based sequential investment strategies. *Mathematical Finance*, 16(2):337–357, 2006.

[13] L. Györfi and D. Schäfer. Nonparametric prediction. *Advances in learning theory: methods, models and applications*, 339:354, 2003.

[14] L. Györfi, A. Urbán, and I. Vajda. Kernel-based semi-log-optimal empirical portfolio selection strategies. *International Journal of Theoretical and Applied Finance*, 10(03):505–516, 2007.

[15] Y. Kalnishkan and M. Vyugin. The weak aggregating algorithm and weak mixability. In *International Conference on Computational Learning Theory*, pages 188–203. Springer, 2005.

[16] B. Li and S.C.H. Hoi. Online portfolio selection: A survey. *ACM Computing Surveys (CSUR)*, 46(3):35, 2014.

[17] B. Li, S.C.H Hoi, and V. Gopalkrishnan. Corn: Correlation-driven nonparametric learning approach for portfolio selection. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 2(3):21, 2011.

[18] Y. Lou, Y. Hong, L. Xie, G. Shi, and K. Johansson. Nash equilibrium computation in subnetwork zero-sum games with switching communications. *IEEE Transactions on Automatic Control*, 61(10):2920–2935, 2016.

[19] M. Mahdavi, R. Jin, and T. Yang. Trading regret for efficiency: online convex optimization with long term constraints. *Journal of Machine Learning Research*, 13(Sep):2503–2528, 2012.

[20] M. Mahdavi, T. Yang, and R. Jin. Stochastic convex optimization with multiple objectives. In *Advances in Neural Information Processing Systems*, pages 1115–1123, 2013.

[21] A. Nowak. Measurable selection theorems for minimax stochastic optimization problems. *SIAM Journal on Control and Optimization*, 23(3):466–476, 1985.

[22] P. Rigollet and X. Tong. Neyman-pearson classification, convexity and stochastic constraints. *Journal of Machine Learning Research*, 12(Oct):2831–2855, 2011.

[23] W. Stout. Almost sure convergence, vol. 24 of probability and mathematical statistics, 1974.

[24] V. Vovk. Competing with stationary prediction strategies. In *International Conference on Computational Learning Theory*, pages 439–453. Springer, 2007.